

REMARKS

This Amendment is filed in response to the Office Action mailed on September 7, 2004. All objections and rejections are respectfully traversed.

Claims 1-87 are in the case.

Claims 43-87 were added to better claim the invention.

No claims were amended.

At Paragraph 2 of the Office Action claims 1-41 were rejected under 35 U. S.C. § 103(a) as being unpatentable⁴ over Han et al. U. S. Patent No. 6,158,017 issued December 5, 2000.

The present invention, as set forth by representative claim 1, comprises in part:

1. A method for enabling recovery from two or fewer concurrent failures of storage devices in a storage array, the method comprising the steps of:
providing the array with a predetermined number of storage devices, including a plurality of first devices configured to store data and

row parity, and one diagonal parity device configured to store diagonal parity, wherein the predetermined number of storage devices n is $p+1$ and wherein p is a prime number;
dividing each device into blocks;
organizing the blocks into stripes that contain a same number of blocks in each device, wherein each stripe comprises $n-2$ rows of blocks;
defining the diagonal parity along diagonal parity sets that span the first devices, wherein the diagonal parity sets wrap around within a group of $n-2$ rows so that all blocks belonging to diagonal parity sets of a stripe are stored in the stripe; and
computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

Han discloses two different dual parity schemes, DH and DH2. Both distribute diagonal parity across the array, by having a separate row of blocks to hold diagonal parity. Scheme DH distributes row parities in a manner similar to RAID5 across rows of the blocks, leaving the rest of the blocks to hold data. DH2 is similar, but moves the row parities to an additional disk added to the end of the array, then leaves the block in the diagonal parity stripe of that disk as a spare (it is outside of any parity calculations). Also, Han stores his diagonal parity results in a stripe across a plurality of his disks (Figs. 5, 6, 7, 8, 9, 10, 12, 13, 14, 16, 17, 18).

Applicant respectfully urges that Han is silent concerning Applicant's claimed novel *providing the array with a predetermined number of storage devices, including a plurality of first devices configured to store data and row parity, and one diagonal parity device configured to store diagonal parity, . . .*

computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

That is, Han discloses storing his diagonal parity blocks on all of his data disks, as shown in his Figs. 3 and 4 where they are stored diagonally, and as shown in his Figs. 5 and 6 where they are stored in a stripe across all of his disks. In sharp contrast. Applicant claims *one diagonal parity device configured to store diagonal parity, . . .*

computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

Further, Han teaches away from Applicant's claimed novel invention by showing two ways to store diagonal parity, both of which spread the diagonal parity out as blocks on all of the disks. The advantage of Applicant's claimed novel invention would be lost on a person of ordinary skill in the art of computer disk storage of data who followed the Han disclosure, where that advantage is set out in the Specification at Page 9 lines 3-14, as:

Advantageously, the inventive technique minimizes the computation load to compute parity stored in the array during failure-free operation. The technique also minimizes the overhead of parity computation, and requires less computation compared to conventional schemes, such as EVENODD, for a given number of data disks. The inventive R-D parity technique speeds recovery by minimizing the total amount of computation that must

be done to recover missing data. In addition, since the invention may be implemented using a concentrated parity technique, where row parity blocks are all stored on the same disk, data disks may be incrementally added to the array without reformatting or recalculating the existing parity information. The only restriction to the addition of data disks to the array is that the maximum number of disks that can be used in the array must be decided beforehand (*a priori*). This restriction is due to the use of diagonals and their lengths depending upon the stripe depth.

Further, Applicant respectfully urges that these advantages would be totally missed by a person of ordinary skill in the art who followed the disclosures of Han.

Accordingly, Applicant respectfully urges that Han is legally precluded from rendering the present invention obvious because of the absence from Han of Applicant's claimed novel *providing the array with a predetermined number of storage devices, including a plurality of first devices configured to store data and row parity, and one diagonal parity device configured to store diagonal parity, . . .*

computing and storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

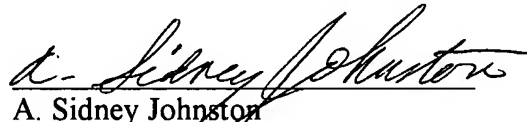
All independent claims are believed to be in condition for allowance.

All dependent claims are dependent from independent claims which are believed to be in condition for allowance. Accordingly, all dependent claims are believed to be in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

A handwritten signature in cursive script, appearing to read "A. Sidney Johnston", is written over a horizontal line.

A. Sidney Johnston
Reg. No. 29,548
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500